Institute for Scientific Computing Research

Student Internship Research Summaries

Center for Applied Scientific Computing

Domain Decomposition Models for Parallel Monte Carlo Transport

Henry J. Alme University of California, Davis

Research Summary

Monte Carlo methods are a desirable way to solve transport problems. Monte Carlo transport methods often provide the best answer for a particular problem, and—in some cases—are the only method available. The objection to Monte Carlo is the large computation time often required for the many samples that must be taken.

Parallel computers have great promise in increasing the performance of scientific computing, but applying parallelism to Monte Carlo transport is problematic. Monte Carlo has unique features that make it more difficult to parallelize when domain decomposition is required (by memory constraints, for example). One such feature is that the computational work for a Monte Carlo problem is localized; it is confined to the parts of the mesh where the sample particles actually travel. This makes conventional parallelization methods less useful.

This work proposes a strategy for parallelizing Monte Carlo transport problems that combines full replication (the most desirable strategy) with domain decomposition, (often required by memory constraints in state-of-the-art problems in scientific computation). The strategy uses various heuristics to estimate which areas of the computational mesh will require the most work, allowing those areas to be replicated more often than areas with less activity. This strategy was applied to several problems, some of them small test problems and some on the scale of real problems. While it was problem-dependent, the scaling behavior, the incremental performance improvement gained by adding more processors to a calculation, of the strategy presented in this work was generally superior to conventional domain decomposition, running two to three time faster in some cases. It required less communication and had better load-balance characteristics.

A Least-Squares Finite Element Method for 3D Neutron Transport Problems

Travis M. Austin University of Colorado, Boulder

Research Summary

Application of a least-squares technique to the 3D neutron transport equation is an exciting alternative to current computational methods. Our method relies on a properly scaled least-squares functional to develop a minimization principle, which is lacking in a standard formulation. The scaling is designed to capture the correct behavior of the solution in different parameter regimes. The minimization principle provides a trusted framework for building the components of an optimal multilevel solution procedure.

Using a finite expansion of spherical harmonics to represent the angular variable leads to a self-adjoint, positive definite system of moment equations. A conjugate gradient method with a block Jacobi preconditioner, that employs multigrid to invert the diagonal elements, is used to approximately invert this system of equations. For nearly all parameter regimes and all diagonal elements, standard multigrid is sufficient as a preconditioner. For certain parameter regimes, however, multigrid performs poorly on the diagonal elements associated with the first-order moments. A new multigrid algorithm has been developed that simultaneously address the three first-order moments.

Our goal is to develop a multigrid routine that has convergence factors independent of discretization size. Smoothing works only on eigenvectors at the high end of the spectrum. Usually, these eigenvectors are exactly the oscillatory errors. With our operator, the other oscillatory errors at low end of spectrum are left unattended. When the smoothing does not truly smooth, the coarse-grid correction becomes worthless as a component in the multigrid algorithm.

For the first part of my summer, I put the finishing touches on code that addressed this previously defined system of equations in 2D. The idea of the algorithm is to augment the standard multigrid algorithm with a special smoothing step. For a two-grid algorithm, we typically have a smoothing followed by a coarse-grid correction and then another smoothing. The new algorithm adds relaxation over divergence-free errors with each step of smoothing. This will “catch” the oscillatory components hiding at the low end of the spectrum. The algorithm we have developed results in convergence factors independent of discretization parameters and bounded by 0.6, which is very good reduction for such a complicated system.

The second half of my summer was spent building the components of the code to solve the full 3D neutron transport problem formulated with least-squares principles. Because of the complication of the new algorithm, particularly with the new finite elements we must introduce, issues of matrix storage, message passing, and effiency of the algorithm arise. Having settled matrix storage and message passing issues, we have begun to develop the necessary components to extend my 2D algorithm to 3D as a preconditioner in the transport code, which is an ongoing project.

Symmetric Wavelets on Arbitrary Topology

Martin Bertram University of California, Davis

Research Summary

Wavelets are used in lossy and lossless compression schemes for scientific data obtained from flow field simulations on supercomputers. Compression is necessary to store, transmit and visualize massive-volume datasets. Most compression schemes are restricted to data defined on regular, rectilinear grids. For compact representation of isosurfaces and material boundaries of arbitrary topology, however, wavelet compression schemes need to be generalized to handle data defined on polygonal meshes.

We constructed new wavelet basis functions defined on subdivision surfaces that have a polygonal base mesh and generate subgrids of regular topology, such as Catmull–Clark subdivision surfaces. Compared to most wavelet approaches on arbitrary topology, our wavelets generalize tensor product basis functions. As a consequence, we obtain B-spline representations for reconstructed data, except in the neighborhood of a few extraordinary points in the base mesh.

We implemented two new wavelet transforms that generalize bilinear and bicubic B-splines to mesh domains of arbitrary topology. To obtain compression, we encode the wavelet coefficients that are sparse or have small absolute values, using a coding scheme that we developed earlier. The overall compression algorithm is highly efficient and provides surface reconstructions at multiple levels of accuracy.

We will use the new surface compression scheme for efficient extraction and visualization of isosurfaces. We also want to develop similar techniques for compression of volumetric data defined on irregular polyhedral meshes.

Material Interface Reconstruction

Kathleen S. Bonnell University of California, Davis

Research Summary

Material interface reconstruction from volume fractions has been approached mainly through the discipline of computational fluid dynamics. We want a generalized visualization solution that can be used for time-varying data sets.

From a rectilinear data set containing *m* different materials, we want to reconstruct the boundary surface of each material. The problem is under-constrained: For general arrangements of volume fractions there is an infinite set of boundary surfaces that are consistent with the fractions. We know that each cell contains a numeric tuple that represents the volume fractions of each material in the cell, and whose sum is 1. Our goal is a simple and robust solution to visualizing the interface.

The general approach is to march through the data in a manner similar to marching cubes, calculating where the interface intersects cell edges, then triangulating these intersection points. Unlike the marching cubes approach, however, which relies on single scalar values for cell vertices, we have had to design an approach that will work with the volume fractions of our data. From our cell-centered representation of the data, we want to create a vertex-centered representation in order to utilize established methods of surface reconstruction where possible. By taking the center of each cell and assigning that cell's volume fractions to the vertex, we can work with the dual of the data set. We then subdivide each dual cell into 6 tetrahedra to reduce the number of cases to be considered for triangulation purposes. From here, we march through each tet, calculating intersection points of the material interfaces. This is accomplished by mapping the volume fractions of the tet vertices into barycentric space. From barycentric space, we calculate material boundary intersections with edges and faces. Finally, we triangulate for each type of boundary: e.g., if we have three materials, a, b, and c, then we triangulate the interfaces between a and b, b and c, and c and a. We expect to get a continuous interface, with coherence over time (when data is coherent over time).

The original approach was proposed by Ken Joy for 2D. I spent the summer working on the 3D solution, searching for a simple enumeration scheme for the different cases (combinations of tet edge intersections and face intersections). We want to create a look-up table similar to marching cubes, one for which we know that no cases have been overlooked and one that can, we hope, be generalized for more than three materials. I have used small datasets to test the discovered cases, and the orientation of the triangulations. This portion is still a work in progress, although only a few more cases remain to be enumerated.

The ASCI visualization team will integrate our work into the terascale browser. I will continue to enumerate the triangulation cases for the three-material solution, solve the problem for more materials, and generalize the current method. I will then compare resulting volume fractions with original data for error analysis, and attempt to solve the problem for unstructured data and for an arbitrary number of materials.

Implementation of an Operator-Split Preconditioner

Thomas A. Brunner University of Michigan

Research Summary

We would like to solve a system of differential–algebraic equations derived from a finite element discretization of the neutron transport equation. When solving this system, a preconditioned GMRES method is used. We need some way to approximate the iteration matrix as a preconditioner. The conventional preconditioner ignores the scattering. For some problems, the scattering can be very important and should be included somehow. The iteration matrix can be inverted in several stages, each stage dealing with a small part of the physics. This operator-split method was implemented in a time-dependent neutral particle transport code.

In order to test the new preconditioner, several simple test problems were devised. In all cases, a small cube with a side of length 1.0 was split up into four zones. Simple one-group isotropic scattering was tested first. The amount of absorption was slowly turned off. A parameter varying the diffusiveness of the problem was varied. The low-diffusion case was difficult to solve, and both the original and the operator-split preconditioner had many linear convergence failures. A diffusion-based preconditioner was tried. Such a preconditioner often, but not always, outperforms the original. A five-group problem with down scattering was also tested. Here, the operator-split preconditioner performs well. Though it converged more rapidly, it also consumed more CPU time than the original.

The operator-split preconditioner is derived with the assumption of small time steps. In the limit where the time steps become vanishingly small, the scattering correction disappears and we recover the original preconditioner. It would be possible to use the original preconditioner for small time steps and to switch to the new one adaptive for longer steps, The limit of large time steps, however, is difficult for both preconditioners, meaning that further work is required for steady-state problems.

Coloring Algorithms for Element-Interpolation-Based Algebraic Multigrid

Tim Chartier University of Colorado

Research Summary

Modern simulations often involve the solution of problems with billions of variables, mandating the use and development of scalable algorithms. While multigrid methods are scalable for many regular grid problems, they can be difficult to develop for the large unstructured grids that many simulations require. Algebraic Multigrid (AMG) attempts to overcome this difficulty by abstracting multigrid principles to an algebraic level so that the algorithm is more automatic and robust. Many researchers see AMG as one of the most promising methods for solving large-scale problems.

My research at Lawrence Livermore National Lab applied AMG to problems of elasticity with complex geometry. This work applies to energy research and the area of nuclear weapons stockpile stewardship. To protect our environment, we must understand how these materials behave in storage over long periods. Acquiring such knowledge is part of the mission of the Department of Energy's ambitious ASCI project.

Lawrence Livermore National Lab researchers are studying AMG in the context of the DOE ASCI project. My research plays a role in their efforts to create scalable algorithms for energy problems. In particular, I am researching coloring algorithms for AMG based on element interpolation (AMGe).

This summer, I continued research on exploiting features available in AMGe to select coarse grids to allow effective interpolation using element interpolation. I continued discussion with LLNL researchers who are also investigating and testing ideas in this area. I developed a very effective 2-level algorithm. Our current research focuses on developing a multilevel algorithm that maintains low element complexity, grid complexity, and convergence factors. It is difficult to balance these competing factors; however, current research shows promise.

Curvature-Based Adaptive Mesh Refinement for Level Set Tracking

Paul Covello University of California, Davis

Research Summary

Level set tracking is used to compute burn tables to be used for detonation shock dynamics code. The level set represents the detonation shock front, which is assumed to propagate like an optical wavefront and thereby follows the eikonal equation. The advantage of using burn tables is that they greatly speed up and simplify the detonation shock dynamics code.

Algorithms exist to track level sets. When given a grid, an initial level set, and speed of detonation shock propagation these numerical codes can predict how the level set will propagate through the rest of the given domain. The advantage of these methods is that they are fast and they work on unstructured grids. The disadvantage is that when the curvature of the level set is high, the accuracy of these algorithms goes down. The way to remedy this is to have higher resolution, or more discrete points, where the curvature of the level set is high.

I am developing an algorithm that keeps a watch on the curvature of the level set: Where the curvature exceeds a certain threshold it will refine that portion of the grid. This new algorithm will be an add-on to the current level set tracking algorithms and also keep the advantage of computational speed and independence of grid type.

Thus far, I have achieved great improvement in accuracy via curvature-based adaptive mesh refinement for non-orthogonal structured grids in two dimensions. My next goal is complete unstructured grids in two dimensions and then do the same in three dimensions.

Large Eddy Simulation of Rayleigh–Taylor Instability Using the Arbitrary Lagrangian–Eulerian Method

Rebecca Mattson Darlington University of California, Davis

Research Summary

This research addresses the application of a large eddy simulation (LES) to Arbitrary Lagrangian Eulerian (ALE) simulations of Rayleigh–Taylor instability. First, ALE simulations of simplified Rayleigh–Taylor instability are studied. The advantages of ALE over Eulerian simulations are shown. Next, the behavior of the LES is examined in a more complicated ALE simulation of Rayleigh–Taylor instability. The effects of eddy viscosity and stochastic backscatter are examined. The LES is also coupled with ALE to increase grid resolution in areas where it is needed. Finally, the methods studied above are applied to two sets of experimental simulations. In these simulations, ALE allows the mesh to follow expanding experimental targets, while LES can be used to mimic the effect of unresolved instability modes.

Classification of FIRST Data Using Support Vector Machines

Matt Giamporcaro Boston University

Research Summary

With the proliferation of information-gathering and -processing resources, scientists have become flooded with more data than they can manage using traditional techniques. The Center for Applied Scientific Computing is developing new algorithms for the exploration and analysis of these large, complex data sets. By applying and extending ideas from data mining and pattern recognition, researchers on Project Sapphire are assembling a toolbox of computational techniques to help scientists extract useful information from these data.

The Support Vector Machine (SVM) is an algorithm that has recently received considerable attention in the pattern recognition literature. It has the potential to be a powerful tool for solving data classification and regression problems. This project has consisted of researching the theory behind SVMs and evaluating how they might be applied to the classification of large data sets. The test data set used has been the FIRST (Faint Images of the Radio Sky at Twenty-cm) set of radio-astronomical data.

The code for a SVM was written in C++. The specific training algorithm implemented is based on an approach presented by Platt (1998) called Sequential Minimal Optimization. The SVM was applied to a subset of the FIRST data set and used to determine whether galactic radio sources might belong to an interesting class known as bent doubles.

The SVM was trained on the data, and test results were compared to those of a decision-tree algorithm implemented using an off-the-shelf software package (C5.0). In training, the performance of the SVM was similar to that of C5.0, in both run-time and cross-validation error rate. During testing on unknown data, there were significant differences in classification performance between the two systems. These results have not yet been analyzed to determine whether they represent better or worse performance on the part of the SVM.

Through the upcoming year, project researchers will continue to evaluate the feasibility of applying SVMs, as well as related techniques such as artificial neural networks, to this classification problem. Directions for continued research on SVMs include: determining whether the SVM performed better than C5.0; exploring the effect of parametric changes to the initial algorithm; and extending the published SVM algorithm to “less crisp” training data.

Explorations of a Stationary Diffusion Equation

Ana Iontcheva University of California, Davis

Research Summary

I worked on three main problems during my LLNL research visit. First, I solved a stationary diffusion equation with the diffusion coefficient represented as a two-dimensional function and also as a matrix. I employed a preconditioned conjugate gradient method with the multigrid V-cycle as a preconditioner, and Gauss–Seidel as the pre-smoother and the post-smoother (forward and backward, respectively).

The second problem was to approach this same diffusion problem, using the classical hierarchical basis.

The third problem was to apply to the solution of this problem the Approximate Wavelet-Modified Hierarchical Basis presented in the paper "Stabilizing the Hierarchical Basis by Approximate Wavelets" by Panayot Vassilevski and Junping Wang. I finished the implementation (in C++) for these three methods, created visualizations, and am comparing their effectiveness.

A GUI Configure System for the KULL Code Packages

Mark Jeunnette University of Chicago

Research Summary

The KULL code is a group of packages, written in FORTRAN, C, and C++, which together work to produce computational simulations of experiments to be carried out at the National Ignition Facility (NIF). Currently, the packages are brought together and compiled using an old and slightly convoluted configure/make system. This is caused mainly by the fact that the packages were brought together under the KULL name after being developed for their own purposes and with their own make systems. The need was seen for a more unified and easier-to-use make system for the entire set of KULL packages.

It was decided that an interactive graphical user interface (GUI) make system would provide the greatest benefit to developers, and ideas were presented as to features to be included in the new system (e.g., cross-checking of package selection, inclusion of library, optimization, and compile flag options).

We began by producing a piece of code that performed the "bare-bones" tasks of allowing the user to set the options needed in a graphical environment and to integrate that with the current configure and make system. That code will be modified and augmented in order to create the final product and allow for additional features later in time. The first program was completed, integrated with the current system, and is ready for use.

Current developers of the KULL code will take over the modification and development of this program, working with it and changing it to fit their needs and the needs of KULL users.

Wavelets in Pattern Recognition

Rachel Karchin University of California, Santa Cruz

Research Summary

I spent the summer researching the use wavelets to decompose images to extract features for pattern recognition. The results are contained in the survey paper "Classifying Images with Wavelet-based Feature Extraction," which will appear as a UC-technical report. The paper includes an introduction to wavelet analysis, wavelet-based classifying systems, and describes in some detail the techniques used to implement such systems, a variety of applications, and results. It contains many pictorial examples to make the concepts quickly accessible, and pointers to further information on the subject.

A number of the techniques described in this paper are new. I added them to the Image Processing Requirements list that the Project Sapphire group is using as a basis for software design.

I also wrote Matlab scripts to experiment with several wavelet-based feature extraction techniques. These include calculating the energy, entropy, high frequency/low frequency energy ratio, and autocorrelation of different levels of a wavelet decomposition of various signals, and producing figures of 1D and 2D images at different resolution levels. The results appear in the survey paper.

Preconditioners for Non-Matching Grids in 3D Elasticity

Chisup Kim Texas A&M University

Research Summary

My research is on the construction of preconditioners in the case where grids on subdomains have been coarsened independently. This coarsening strategy is a common practice in algebraic multigrid applications, which is one of the main projects at CASC, and leads to subdomain interfaces with non-matching grids. Then, to retain the approximation properties on these interfaces, the mortar finite element method, a non-conforming finite element method, is applied.

In the 3D linear elasticity application (away from the incompressible limit), such an interface is a 2D domain, where the mortar finite element method has not been thoroughly analyzed and applied. Joseph Pasciak, Raytcho Lazarov, who were visiting here from Texas A&M University, Panayot Vassilevski at CASC, and I developed a stable and efficient algorithm that can handle such 2D interfaces. A code for this algorithm is being developed and this result is expected to be published in an academic journal.

Data Mining in Astronomical Data

Imelda Kirby University of Washington

Research Summary

Data collection threatens to overwhelm our capacity to organize and review it for scientific information and inference. Project Sapphire in the Center for Applied Scientific Computing at LLNL is developing new techniques for analyzing massive quantities of data to extract the features that are useful for scientists engaged in the research.

Our project in data mining for astronomical data began with the installation of IRAF, and other software for astronomical images. I learned how to handle the MACHO (Massive Compact Halo Object) images, and to research different automated ways of detecting asteroids in the terabytes of data MACHO contains. I also researched Self-Organizing Map (SOM) algorithms to see if they could be applied to extract features from large astronomical databases.

Software Infrastructure for Large-Scale Partial Differential Equation Simulation Codes

Matthew Knepley Purdue University

Research Summary

My work focused on providing software infrastructure for large-scale partial differential equation (PDE) simulation codes. I become involved in the Equation Solvers Interface (ESI) effort to develop an interface standard for linear and nonlinear sovlers, which led to a report co-authored with Andy Cleary. This report outlines a proposal for a framework based upon an Operator interface, which encapsulates the behavior of an arbitrary finite dimensional operator.

I have also prepared a draft proposal for Frameworks, which examines the object-oriented design of simulation software for mesh-based PDEs, as well as portability and interoperability issues. Furthermore, I looked at solving the nonlinear system that arises from fluid flow problems, using a novel projection scheme related to traditional Krylov methods.

Toward a Parallel Three-Dimensional Ideal Magnetohydrodynamics Solver

Joseph Koning University of California, Davis

Research Summary

The simulation of ideal magnetohydrodynamics (MHD) requires the solution of a fluid dynamics problem coupled with electrodynamics. An interesting problem in this field is a compact toroid moving through a plasma contained within a tokamak fusion device. The idealization of this problem is a superconducting sphere moving through an ideal plasma. The sphere's motion will excite waves in the plasma. The energy loss through the interaction of the sphere with the plasma and magnetic field will result in the sphere slowing and eventually stopping. This problem has been treated analytically by Newcomb. The superconducting sphere in an ideal plasma serves as a proof-of-concept for the problem of a compact toroid interacting with a tokamak plasma. This problem is three-dimensional and contains fluid phenomena such as shocks and electrodynamic phenomena such as Alfven and magnetosonic waves.

I will use an Arbitrary Lagrangian Eulerian (ALE) method couple; the ALE method to treat the fluid dynamics of the plasma; and the VFE method to treat the electric and magnetic fields. The ALE method has the ability to track interfaces; in the simulation, the sphere moves through the plasma making front tracking necessary. The VFE was chosen for its ability to treat the electric and magnetic field's continuity as it is physically. Vector basis functions maintain normal continuity for the magnetic field and tangential continuity for the electric field.

I have been researching the abilities of the ALE methods available for use as a basis for the MHD simulation. The fluid dynamics framework I chose is KULL, an object-oriented program with an interpreted interface with the ability to treat a fluid using pure Lagrangian, free Lagrangian, Eulerian or ALE. Currently KULL does not have finite element capability. In the latter part of this year I developed an object-oriented VFE framework which can create the Hilbert and Gramm matrices for the four types of basis functions (cell, node, edge and face) as well as mapping matrices for improved performance. This framework will then be incorporated with KULL to complete the framework for the MHD simulation. I have also researched parallel linear algebra libraries and have become proficient in the use of PETSc, the portable extensible toolkit for scientific computing.

In the future, I will be coupling the VFE framework and KULL. This will entail new methods for manipulating the vector basis functions and advecting the magnetic field. Once I complete the proof-of-concept simulation, the simulation of the real physical system will require further extension of the VFE and KULL frameworks.

Superresolution of Buried Objects in Layered Media by Near-Field Electromagnetic Imaging

Sean K. Lehman University of California, Davis

Research Summary

In many tomographic wave (either electromagnetic or acoustic) imaging systems, wavelengths and distances are such that the object under evaluation lies in the near-field. Near-field scattering includes non-radiating or evanescent fields. Evanescent fields carry high spatial frequency information. Thus, if they can be incorporated into tomographic reconstruction algorithms, higher spatial resolutions can be achieved than those predicted by the classical (Rayleigh) resolution limit. This is known as "superresolution." Most current tomographic reconstruction methods neglect the evanescent fields.

I have been developing a physics-based theory (as opposed to an engineering-motivated approach) that seeks to understand how the evanescent fields are scattered and transformed when measured. Analysis shows that the spectrum of the evanescent fields shifts into propagating regions during the scattering and measurement process. This has never been considered previously. This "smearing" of spectral energy would be considered as noise in the current techniques which do not account for the evanescent fields.

My analysis is in a planar geometry. That is, a planar measurement system. After the dissertation, I would like to develop it for an annular or circular geometry.

Benchmarks and Models for Time-Dependent Grey Radiation Transport with Material Temperature in Binary Stochastic Media

David Miller University of California, Davis

Research Summary

We present benchmark calculations for radiation transport coupled to a material temperature equation in binary random media. The mixing statistics are taken to be homogeneous Markov statistics where the material chunk sizes are described by Poisson distribution functions. The material opacities are taken to be constant. Benchmark values for time evolution of the ensemble average values of material temperature, energy density, and radiation transmission are computed via a Monte Carlo type method. These benchmarks are used as a basis for comparison with three other approximate methods of solution.

One of these approximate methods is simple atomic mix, which is seen to consistently over-absorb resulting in lower steady state radiation transmission and material temperature. The second approximate model is an adaptation of what is commonly called the Levermore–Pomraning model, which we refer to as the standard model. It is shown to consistently under-absorb resulting in higher steady-state radiation transmission material temperature. We show that recasting the temperature coupling as a type of effective scattering can be useful in formulating the third approximate model, an adaptation of a model due to Su and Pomraning that attempts to account for the effects of scattering in a stochastic context. We show this last adaptation demonstates consistent improvement over both the atomic mix and standard models.

Non-Uniform Load Balancing in SAMRAI

Scott Morris University of Utah

Research Summary

The end product of the CSAFE project at the University of Utah is a PSE (Problem Solving Environment) that couples MPM (Material Point Method) simulating container dynamics with a continuum firespread code. Both of these pieces sit upon the SAMRAI adaptive mesh refinement infrastructure. The incorporation of the MPM code with SAMRAI required the creation of a new particle data type. These particles are free to move about the mesh and as such non-uniform distributions of the computational workload result across the mesh. Since the CSAFE project is part of the ASCI initiative, the goal is to achieve a PSE that is scalable to large numbers of processors. To facilitate this goal, SAMRAI requires a non-uniform load balancer. The objective of my work was to research an algorithm for non-uniform load balancing within an AMR framework and to implement a general purpose load balancer in SAMRAI.

The traditional approach to dividing work among processors on a grid calculation is to use recursive bisection. This approach is useful when given a large, single domain. With adaptive mesh refinement, however, we are given the task of balancing multiple levels of refinement and irregular domains. Instead, we use a modified form of recursive bisection that can partition a collection of boxes into any number of boxes of a specified size and produce a smaller leftover box as well. This is a heuristic for chopping boxes that prepares for the next step of the algorithm: a bin-packing stage in which chopped boxes are assigned to processors. In practice, this heuristic has proved to have merit. For the linear advection sample SAMRAI code we have achieved average estimated load balances in the 90–100% range.

Further testing of the algorithm is required as different applications using SAMRAI have different needs. Specifically we will be testing the non-uniform load balancer with the MPM code currently under development at the Utah CSAFE site. Other refinements to the algorithm include consideration of communication costs and a more efficient bin-packing algorithm.

Parallel Sparse Matrix–Vector Multiplication

Jason Perry University of Kentucky

Research Summary

Multiplying a sparse matrix by a vector is a basic task in nearly every scientific simulation code run on parallel computers. My task was to attempt to speed up this operation by adding threads to an already parallel code. The proposition rests on the dual hardware structure found in supercomputers and clusters of networked workstations. Each node of the machine (or cluster) is a complete separate computer, and messages can be exchanged between nodes to solve larger problems. Within each of these nodes is also contained a number of processors that share the memory within the node. Preexisting codes perform matrix–vector multiply in parallel by essentially dividing the rows of the matrix up among the nodes. These codes use only message passing, and therefore either ignore all but one processor within each node or pack multiple processes within one node, which causes unneeded communication overhead. Multiple threads of execution within a shared memory space allow work to be divided between processors in a node, without the need of any communication within the node itself. A program code that combines message passing and threading should reduce the amount of inter-process communication on parallel computers, thereby speeding up the computation.

My first step in this experiment was to write code that would allow me to input matrices for testing. I wrote a program in C to read matrices stored in the standard Harwell–Boeing format. I then proceeded to write from scratch a parallel matrix–vector multiplication code that uses the MPI message passing standard and no threads. This code was tested on both the IBM Pacific Blue supercomputer and a cluster of DEC Alpha workstations. I tested and timed this code and refined it until I felt it was sufficiently fast, as well as representative of existing codes. Then, I added code that uses the Pthreads threading standard to thread the multiplication within each node. I created several new versions of the program, each with a different method of threading the loop that does the multiplication. Some methods use dynamic load balancing to evenly distribute work among the threads, while others simply divide up the rows of the matrix beforehand. Some versions perform the multiplication and communication between nodes at the same time, by assigning different tasks to different threads. These were all extensively tested with a variety of matrices from scientific applications, as well as specially generated regular matrices with up to 300 million entries. The goal was that my threaded code would run faster than a non-threaded code running the same problem on the same number of processors.

The results have been positive, showing that threading can provide a modest speedup if done correctly. Also, important insights have been gained into some of the performance limitations of new parallel machines. I have prepared extensive documentation, which will be further revised, and which I hope to publish, and have given two presentations on this work.

Scalable Domain Decomposition Algorithms for Resolving Contact Surfaces in ALE Computations

Tim Pierce University of California, Davis

Research Summary

In the Lagrangian approach to modeling material media, the overall volume of material under consideration is subdivided into a discrete set of zones, together with their connecting faces, edges, and nodes: i.e., a mesh. The connectivity of the mesh determines which nodes belong to each zone, from which can be deduced which zones contain a given node. This information establishes the proximity relations of all zones.

As the material moves in time under the influence of internal stresses, body forces, and boundary conditions, the mesh follows the material, so that a given zone always contains the same material. For an explicit time advancement scheme, the behavior at a given node or zone over a single timestep is dictated only by material in its own and neighboring zones, which can be determined from the mesh connectivity. The timestep is chosen, with reference to the speed at which signals can propagate through the material (typically the sound speed), so that this locality of influence is a valid assumption.

If the model includes more than one discrete body, the interaction of the bodies must also be considered, and mesh connectivity is no longer sufficient to determine proximity. The bodies may come together or separate, or may slide along each other, each exerting a boundary force on the other. Though the connectivity of the mesh representing each individual body is constant, which zones are adjacent across the contact surface is entirely dynamic and unpredictable.

A standard approach to parallelizing a Lagrangian dynamics code involves dividing the mesh into a number of submeshes, or subdomains, of approximately equal size, and assigning the subdomains to separate processors. A good decomposition minimizes communication, as only neighbor information is required. Furthermore, since the connectivity is constant in time the communication pattern can be set up during an initialization phase, and only the variable data itself updated each cycle.

This simple domain decomposition method of parallelization breaks down at contact surfaces, or at least an additional mechanism must be supplied to a) determine proximity relations across contact surfaces, b) distribute the calculation of the contact forces across the parallel machine, and c) gather the necessary data from across the machine together on the processor where that part of the surface is to be calculated.

The task of calculating contact forces can be divided into two steps, contact detection and contact enforcement. First, one must determine if and where two surfaces are in contact. Then, one can apply suitable balancing forces at those points.

In a serial environment, the detection problem can always be solved by brute force. The position of each node on one side can be compared to the positions of all faces on the other side, the closest face determined, and then penetration checked. In a parallel environment, however, even the brute force approach is not available, since the closest face on the other side may be on a different processor. As time changes, the closest face and perhaps the processor assigned to that face may also change.

As opposed to the relatively simple problem of parallelizing a Lagrangian dynamics code on a single mesh, the parallelization of contact surfaces provides formidable challenges, particularly if the task must be done in a scalable fashion as the mesh size and processor count are increased proportionately.

The purpose of this research is to develop a particular algorithm designed to provide a truly scalable solution, to implement it in a major dynamics code (ALE3D), and to demonstrate scalability well beyond 1,000 processors.

ROAM Using Surface Triangle Clusters

Alex Pomeranz University of California, Davis

Research Summary

As terascale datasets become commonplace, new methods are needed to visualize these enormous amounts of data accurately and quickly. Current display techniques are CPU-heavy and do not come close to utilizing the full potential of the graphics cards in the consumer market. Our accelerated display technique builds upon the view-dependent optimization technique presented in ROAM (Duchaineau) with the intention of allowing the user to view arbitrary 3D surfaces in real-time.

The original ROAM paper used a BinTree hierarchy to store various approximation levels of our 2D mesh/terrain. By taking some cut across this tree, such that a certain number of polygons is displayed or a certain maximum error threshold is met, we can optimally display our mesh given the constraints of our error metric in real time. Graphics cards today are capable of rendering 9 million shaded and textured triangles per second; however, the original ROAM paper sent only 6,000 triangles/frame at 30 frames/second, for a total of 180,000 triangles. This amounts to a 2% utilization of our graphics hardware.

In our new method, we take a more coarse-grained approach. After the cut across the BinTree hierarchy is found, rather than display individual triangles, we display triangle clusters, already precomputed and prepackaged for quick output. Assuming 256 triangles/cluster, we would then be pushing 180,000 triangle clusters at 256 triangles/cluster, or 4.6 million triangles/second. This not only gives a much better approximation of our actual mesh, but it also offers a 25 times increase in the utilization of our graphics hardware.

We intend to adapt this method to 3D meshes. This gives our algorithm a much wider base of potential users. Currently, we are working on the initial prototype for this system. We already have a large-object oriented framework up and running, and are incrementally adding functionality to it. Some issues to be addressed are the adaptation of the algorithm to 3D meshes, “prepackaging” triangle clusters, and “paging clusters to disk” issues.

A Multigrid Strategy for Accelerating Steady-State Computations of Waves Propagating with Curvature-Dependent Speeds

Jonathan Rochez University of California, Davis

Research Summary

The need for a fast and accurate representation of lighting times of detonation of materials has arisen in several problems. Such a representation is called a burn table and is simply an array of times at which the detonation burn front crosses a particular point relative to the detonation point(s). In a high-explosive material, the creation of a burn table will allow the elimination of solving chemical reaction ordinary differential equations (ODEs) and feed in source terms to the reactive flow equations for solution of the system of ignition of the high explosive material. Standard iterative methods show a quick reduction of the residual followed by a slow final convergence to the solution at high iterations. Such systems are excellent choices for the use of multigrid methods to speed up convergence, even on a nonlinear system such as this. During the past year, we have attempted to develop this code in 2D and 3D using two differencing schemes for the multigrid method, and numerical steady-state solutions to the eikonal equation on a rectangular grid were conducted.

Numerically, two different approaches have been investigated with the multigrid method. The first approach is organizing the calculation point by point on the grid using a finite differencing scheme. The advantages of this approach include its speed in calculation time and ease of implementation, while its downfall is having to remain on a regular orthogonal grid. During this past year, code was developed to solve the eikonal equation on a unit square geometry with one or several detonation points involved. Comparisons of several sized problems with several multigrid V-cycles have been compared. In all cases, the speed was significantly faster for the multigrid method versus a regular calculation on a single grid. Also, the larger the problem size became, the faster the speed-up that was observed. The best experimental results as seen in the comparison table below show a factor of 60 reduction in wall clock time.

The second approach is one of using a finite element scheme in a zone by zone organized calculation. Advantages include unstructured grids for the formation of non-regular geometries and fewer data points, while the disadvantages include larger claculation times and being somewhat more difficult to implement. Code has been developed to test this method as well. No results have been obtained yet, due to the unstable and non-converging behavior of using a linear “hat-function” basis set.

Current plans for improvement include expanding the geometries for the finite differencing scheme with the incorporation of curvature in the calculation of the burn table. Error analysis and comparison to physical data must be performed as well. The finite element approach will be fully implemented, while much work is still needed on the grid generation section—specifically prolongation and relaxation between grid levels.

Comparison of single grid versus multigrid for select problems

Problem	Calculation time (sec)
2D, 128x128, single grid	1889.33
2D, 128x128, multigrid	314.29
3D, 33x33, single grid	66723.40
3D, 33x33, multigrid	1023.49

Gamma-Ray Bursts: Numerical Modeling of Electron–Positron Pair Plasmas in Cataclysmic Astrophysical Phenomena

Jay Salmonson University of California, Davis

Research Summary

Despite three decades of intense scientific scrutiny, gamma-ray bursts have remained one of astronomy's biggest unsolved mysteries. Recent observational breakthroughs have allowed us to learn much about these big, brief, brilliant bangs seen from across the cosmos, but their origin remains a mystery. In this work, we study three progenitor models: a neutron star binary system near its last stable orbit, a charged black hole, and the collapse of a globular star cluster. All of these scenarios result in a common theme, the relativistic expansion of a super-heated electron–positron–photon plasma. Thus, we study the evolution of, and emission from, this plasma as it might result from these three progenitors using numerical general relativistic hydrodynamic simulations. This emission is then compared with that of gamma-ray bursts to test the feasibility of each of these models as a gamma-ray burst progenitor.

Implementing the Full Approximation Scheme Algorithm

Kevin Scully University of California, San Diego

Research Summary

We have created an implementation in C of the FAS (Full Approximation Scheme) multigrid algorithm developed by Achi Brandt to apply to discretizations of nonlinear partial differential equations (PDEs). This algorithm may permit the solving of nonlinear PDEs with speed and efficiency comparable to that of the linear multigrid method on linear PDEs. While not much theory exists to support the effectiveness of this algorithm, its power has been observed in certain implementations.

Under the guidance of Carol Woodward, an FAS algorithm using a Point Newton finite difference weighted Jacobi smoother, semicoarsening to generate coarse grids, full-weighting and injection for restriction and linear interpolation, was developed to handle linear and nonlinear problems in up to three dimensions, where the domain is of tensor–product type. This code reproduces results of benchmark Matlab FAS codes written by Van Henson on elliptic linear and nonlinear Dirichlet boundary problems.

This code uses many data structures and functions to manipulate these data structures from the *hypre* library, developed at LLNL. In particular, it makes liberal use of the "hyprebox" data structure and has been written to handle multiple box domain partitions with an eye to a parallel computing implementation. Many of the routines dealing with boxes (such as the restriction and interpolation functions) were patterned after comparable routines in the linear multigrid solver PFMG, also developed here at LLNL by Rob Falgout and others.

I plan to explore the possibility of interfacing this code with finite element discretizations of Einstein's equations produced by Michael Holst's MC package. To use this algorithm in space-time finite element discretizations (where the finite elements have a time component), I will have to extend this algorithm to four dimensions.

Flexible GMRES in CVODE and PVODE

Alexandru Tamasan University of Washington

Research Summary

CVODE is a nonlinear solver for ordinary differential equation initial value problems developed by A. Hindmarsh and S. D. Cohen at LLNL. PVODE is a parallel version. The algorithm executes an outer time loop with an adaptive time step. Newton's method is applied in solving an implicit time discretization scheme (e.g., backward differentiation or Adams) on each time step. A linear solver is required for the corresponding Jacobian equation. An iterative method based on the GMRES algorithm is available, called CVSPGMR. In 1993, Saad noticed that for the price of a little extra storage, GMRES accommodates a variable right preconditioning. This new linear solver is called Flexible GMRES. My task was to implement it within the CVODE package, resulting in a module called CVSPFGMR. Despite the fact the solution is formed differently, the residual solves the same least squares problem as in the GMRES algorithm. Hence we can use the same stopping criteria. Moreover, the convergence results for the Inexact Newton–GMRES carry over to the Inexact Newton–FGMRES algorithm.

In order to compare the two methods, a test problem previously run with CVSPGMR was run with CVSPFGMR. A variable preconditioning on the right was realized by requiring a fixed tolerance on the number of Gauss–Seidel preconditioner iterations, rather than a fixed number of iterations. For similar performance in terms of measured outputs, the numbers of iterations could be reduced by a factor of more than two, making a variable preconditioner is a better choice. The disadvantage is the increase in the memory space.

Work in progress deals with the mixed initial-boundary value problems for a 3D heat equation. PVODE runs with CVSPFGMR as the linear solver, CG with SMG from the *hypre* package as the preconditioner, combining two independently developed CASC codes.

Unsplit ALE Using an Augmented Lagrangian Mesh Velocity

Jay F. Thomas University of California, Davis

Research Summary

The Arbitrary Lagrangian–Eulerian (ALE) method has been successfully applied to the simulation of variety of problems including impact phenomena,material processing (e.g. metalforming) and fluid–structure interaction. The ALE method combines the Lagrangian approach, in which the computational mesh moves with the material, with the Eulerian approach wherein the mesh is fixed and the material flows through it. The Lagrangian approach simplifies application of boundary conditions and tracking of material interfaces, but the mesh can become severely distorted or entangled—ventually halting the computation.

The Eulerian approach allows arbitrarily large deformations and turbulent flow but treatment of boundary conditions and material interfaces may require special accomodation. The ALE formulation attempts to circumvent these difficulties and obtain the advantages of both methods by allowing the mesh to be moved in an arbitrary manner. The issue becomes one of selecting the mesh motion so as to control grid distortion and or solution error.

ALE methods have typically been implemented as uncoupled or “split” procedures consisting of a Lagrangian phase followed by an Eulerian or advection phase in which the mesh velocity is calculated (remeshing) and the state variables are remapped to the updated mesh positions. There is no time step associated with the advection phase. The advection step can add considerable computational expense, particularly if frequent remeshings are needed. Many ALE codes adopt a mesh relaxation procedure to move the nodes with an eye towards controlling grid distortion. Mesh relaxation is based upon elliptic mesh generation methods. The coordinates of the nodes are formulated as the solution to an elliptic problem and Jacobi iteration is used instead of a direct solution method. This results in a stencil that can be applied efficiently over the mesh. Only a few sweeps of Jacobi are used at each time step.

The goal of the effort is to develop and investigate unsplit ALE formulations and associated computational procedures with particular attention to control of mesh distortion. With an unsplit formulation, there is no operator split and hence no separate Lagrangian and Eulerian phases. In principle, greater accuracy is possible in comparison to the split procedure. Controls for mesh distortion are specified directly as part of the differential equation for the mesh velocity, as opposed to indirectly via mesh relaxation. The mesh may thus be considered to deform, and potentially adapt to the solution, continuously. Initial effort is focused on mesh velocities of the form: $u + f(t,t)$ where u is the material velocity, x is spatial position and t is time. f is a "correction" term to prevent mesh distortion and entanglement. Preliminary experiments performed using the inviscid Burgers' equation indicate promise for the method.

Future plans are to develop the formulation for a 1D system of equations to aidin chosing appropriate candidates for the $f(u,x,t)$ term. The technique would then be extended to 2D.

Sparse Multiple Semicoarsened Grids

Heidi K. Thornquist Rice University

Research Summary

Much work has been done developing robust multigrid methods for solving systems arising from finite difference, finite volume or finite element discretizations of the diffusion equation, in the presence of discontinuous coefficients and anisotropy. However, all of these methods have higher computational and memory costs than methods based on point smoothers and simple local grid transfer operators. The goal of this work is to develop an algorithm that is both robust and efficient.

The new method we developed is a modification of the multiple semicoarsened grids (MSG) method of Naik and Van Rosendale. The robustness of this method was studied by Oosterlee and Wessling and a new, more efficient variant of this algorithm was recently proposed by Washio and Oosterlee. Our approach for improving efficiency is based on the relationship between the MSG algorithm, the hierarchical-basis multigrid method, and the sparse grid methods of Zenger. Our new method, denoted SparseMSG, has been tested on a variety of 2D and 3D scalar diffusion problems, where, in iteration count, it beats LLNL's PFMG. In the future, we intend to optimize SparseMSG for parallel implementation.

Selecting Data Mining Tools

Ben Tobin Northern Arizona University

Research Summary

Scientific data that has been carefully collected over time can now be exploited by new data mining technology. For use in anticipated LLNL terabyte applications, a data mining tool must be able to handle hundreds of data columns, handle flat files, and be parallel. To evaluate data mining options for LLNL, I began with the *Two Crows* report, a 370-page report evaluating the 25 largest data-mining products on the market, of which 8 support parallel architectures. In consultation with CASC scientists, we have narrowed the field to five data mining products: Intelligent Miner, SPSS, CART, Darwin, and Clementine.

We presented the results of our trials and the ease of interaction with these various products. It is possible, however, that the data-mining tool that we need has not yet been developed.

Selected references

Herbert Edelstein, *Two Crows Data Mining '99 Technology Report*, Two Crows Corporation, Potomac 1999.

JSL Decision Systems Inc. Clementime
http://www.isldsi.com

SPSS
http://www.spss.com/datamine

Salford Systems, CART
http://www salford-svstems.com

Neo Vista Software, Decision Series
http://www.hp.com/fsi/partners/partner585.html

Silicon Graphics, MineSet
http://www.sgi.com/software/mineset

Thinking Machines Corporation, Darwin
http://www.think.com

Domain Decomposition on Non-Matching Grids

Stanimire Tomov Texas A&M University

Research Summary

I worked with Raytcho Lazarov and Panayot Vassilevski on two articles, “Interior Penalty Discontinuous Approximations of Elliptic Problems” and Discontinuous Approximations of Advection-Reaction Equations. The first article studies a penalty discretization of second-order elliptic boundary value problems, which is first-order accurate if one uses piecewise linear elements. The second one proposes a conservative (finite volume) approximation of advection-reaction problems. I finished the implementation (in C++) for both methods and provided extensive numerical results. I worked on *a posteriori* error estimators and adaptive refinement and applied some local refinement techniques for the penalty formulation for diffusion problems. Work is now in progress is work on mortar approximation for convection-dominated diffusion problems (with upwinding), and derivation of local error estimators for purely convection-reaction problems.

A significant part of my work was on 3D mesh generators, design and implementation in C++ of data structures for multilevel adaptive grid refinement, implementation of various local error estimators, and tests of their effi-ciency. The experience of Mike Holst and Randy Bank, who were visiting CASC at the same time and with whom I had discussions, turned out to be very helpful. For the 3D mesh generation I used NETGEN, a 3D mesh generator based on advancing front method, efficient in generating coarse meshes. The mesh refinement is based on Douglas Arnold’s bisection algorithm.

Interface Design for Visualization Applications Using Adaptive Mesh Refinement

Kevin Vlack University of Illinois, Urbana–Champaign

Research Summary

Adaptive Mesh Refinement (AMR) has recently become a popular method for focusing computational resources on subregions of a domain relevant to the stability and accuracy of computational experiments. Subregions created through AMR may then be re-assembled to represent the entire domain, but the complexity of this reassembly complicates numerical analysis, and visualization tools have become necessary to study the results of these experiments. Such tools require efficient methods of data representation which can be easily interpreted by a human researcher.

AMRVIS, developed by Vince Beckner of Lawrence Berkeley National Laboratory, is an application used to visualize data from files produced by *hypre*, a library for AMR solvers for problems in fluid dynamics and radiation transport. Our new version provides an interface improvement to its previous release, reducing the clutter in the GUI and introducing menu item accelerators for every option in the tool, such as viewing datasets and selecting different levels of refinement for view. A module has been added for viewing older versions of the data files produced by the *hypre* code, which can now be viewed in parallel. A new 1D line plotting utility allows viewing the values along a horizontal of vertical line of the subdivided computational domain on a traditional Cartesian plot. This feature also supports animation and level-switching. Such capabilities are important to the end-user and the *hypre* programmer for purposes of software evaluation and for understanding the physical processes behind experimental results.

The *hypre* code is currently under development for parallel 3D simulations for fluid dynamics and diffusive radiation transport. As the code nears completion, the code for AMRVIS must be able to interactively support the visualization of selections of a 3D dataset. While the current software and its new version allows viewing in 3D, its methods in data selection and visualization are still rather simple and must undergo further development.

Data Locality Optimization for the ROSE C++ Preprocessor

Christian Weiss Technische Universitaet

Research Summary

The ROSE preprocessor is part of LLNL's Overture project for solving systems of partial differential equations on general multipatch grids. I implemented a two-dimensional red–black Gauss–Seidel relaxation method using the A++ class library to get familiar with the usage and implementation of ROSE. ROSE was not able to transform codes which contains multidimensional arrays. Hence, I modified ROSE so that it is now able to correctly handle and transform C++ code which uses A++ arrays with more than one dimension.

To automatically apply data locality transformations with a preprocessor, several features are needed within the code transformation tool. I identified additional features I would need within ROSE to implement data locality optimizations. Array padding was selected, since it requires the least features to be added to the ROSE preprocessor. Array padding is a technique that changes the memory layout of arrays to reduce the number of conflict misses in an application. The library I implemented provides array padding heuristics as well as tools to apply padding to array descriptions. The library can be used to implement a code transformation within ROSE as well as to implement runtime support for A++ array constructors to improve the memory layout of dynamically allocated arrays.

A Reduced Grid Method for a Parallel Global Ocean General Circulation Model

Michael Everett Wickett University of California, Davis

Research Summary

A limitation of many explicit finite-difference global climate models is the timestep restriction caused by the decrease in cell size associated with the convergence of meridians near the poles. A computational grid in which the number of cells in the longitudinal direction is reduced toward high-latitudes, keeping the longitudinal width of the resulting cells as uniform as possible and increasing the allowable timestep, is applied to a three-dimensional primitive equation ocean-climate model.

This "reduced" grid consists of subgrids which interact at interfaces along their northern and southern boundaries, where the resolution changes by a factor of three. Algorithms are developed to extend the finite difference techniques to this interface, focusing on the conservation required to perform long time integrations, while preserving the staggered spatial arrangement of variables and the numerics used on subgrids. The reduced grid eliminates the common alternative of filtering high-frequency modes from the solution at high-latitudes to allow a larger timestep and reduces execution time per model step by roughly 20%. The reduced grid model is implemented for parallel computer architectures with two-dimensional domain decomposition and message passing, with speedup results comparable to those of the original model. Both idealized and realistic model runs are presented to show the effect of the interface numerics on the model solution.

First, a rectangular, mid-latitude, flat-bottomed basin with vertical walls at the boundaries is driven only by surface wind stress to compare three resolutions of the standard grid to reduced grid cases which use various interface conditions. Next, a similar basin with wind stress, heat, and fresh water forcing is used to compare the results of a reduced grid with those of a standard grid result while exercising the full set of model equations. Finally, global model runs, with topography, forcing, and physical parameters similar to those used for ocean-climate studies, are advanced to a near equilibrium state for both the reduced grid and the standard grid. Differences between the two are presented for typical fields of interest, and very little degradation of the solution due to the reduced grid is observed.